



Digital Regulation Cooperation Forum



Transparency in the procurement of algorithmic systems:

Findings from our workshops with vendors and buyers

Note: This discussion paper is intended to foster debate and discussion among our stakeholders. It should not be taken as an indication of current or future policy by any of the member regulators of the Digital Regulation Cooperation Forum (DRCF).

Executive Summary

Algorithmic Systems (AS) are present across almost all areas of our lives and are a key component to many innovations in digital technologies that we see today. AS can help businesses automate tasks, both enabling greater efficiency and unlocking new capabilities for businesses and their users. They can also enhance the provision of public sector services to citizens. However, AS can also pose significant risks if deployed without due care. For example, AS may exhibit harmful biases that could lead to unfair or misleading outcomes for people and businesses.

In 2021, the DRCF established an algorithmic processing team to explore the impact of algorithms across our industries and regulatory remits. As part of this work,¹ we committed to undertaking further research to explore the role regulators might play in promoting transparency in algorithmic procurement.

In the autumn of 2022, we held two workshops with 23 vendors and buyers of AS to explore transparency during the process of procuring AS. A key aim of the workshops was to explore how the procurement of AS works in practice, including what information is shared, how it is shared and who is responsible for sharing it. In these workshops, we explored:

- How procurement of AS takes place.
- Information that vendors shared with buyers (and that buyers shared with vendors).
- Key barriers to transparency regarding key features of AS.
- Potential solutions that can help to overcome these challenges.

This paper discusses the insights and key findings from these two workshops. Our key findings are:

Finding #1: There is not a one-size-fits-all approach to achieving transparency in the procurement of AS.

We found:

- The procurement of algorithms goes beyond acquiring off-the-shelf systems and encompasses acquiring data sets which can be used to train AS, among other types of procurement.
- Different organisations have different processes for procuring AS.
- Information exchanged between vendors and buyers during the procurement of AS can vary from case to case.
- Assessments are undertaken by both vendors and buyers in relation to whether AS will be deployed appropriately and lawfully.
- A variety of support (including no support) is provided by vendors to buyers after procurement has taken place. However, the relationship between a buyer and vendor often needs to continue beyond the initial procurement, for example, to address changing contexts.

¹ Digital Regulation Cooperation Forum (2022) Plan of work for 2022 to 2023. Available: <https://www.gov.uk/government/publications/digital-regulation-cooperation-forum-workplan-2022-to-2023>

Finding #2: Buyers can lack the technical expertise to effectively scrutinise the AS they are procuring, whilst vendors may limit the information they share with buyers.

We found:

- Some buyers lack the technical knowledge to understand performance metrics and adequately scrutinise the AS that they wish to procure.
- There is a concern among some vendors that other competitors can cherry pick performance metrics which can over-state the effectiveness of an AS.
- Vendors can find it difficult to strike the right balance between the desire to provide information about an AS with the need to protect commercially sensitive information.

Finding #3: There are some opportunities for achieving greater transparency across algorithmic procurement emerging.

We found:

- Certification, standards, and guidelines may have a role in setting out agreed metrics that vendors can share with buyers.
- Templates and registers can function as a guide to help vendors understand good practice in relation to the types of information to share with buyers.
- An AI auditing market may help to provide assurance to buyers that a technology has met a certain standard.

This paper reflects our discussions with workshop participants on the specific topics above and is not a complete assessment of transparency in the procurement of AS. We encourage further work by regulators, industry, academics, and civil society to explore and establish good practice.

Transparency in the context of algorithmic processing will continue to be a focus of the DRCF's work, as well as for other regulators, governments and regions who are developing approaches to AI regulation and governance. We hope that this paper will inform conversations and interventions to support greater transparency in the procurement of AS.

Introduction

As part of the DRCF's algorithmic processing work, we have previously published two discussion papers, first on '*The benefits and harms of algorithmic systems*'² to help articulate the nature and severity of algorithmic risks and the measures that are needed to mitigate them, and second on the algorithmic audit landscape in '*Auditing algorithms*.'³

By AS, we mean a statistical model that processes data to produce an output or to make a decision. These models can be developed using machine learning methods or simpler statistical techniques. Individual models are often combined with others to create a single 'algorithmic system' (AS) that performs a given function.

We understand that AS are being used in many ways across the economy, for example:

- Supermarkets may use algorithmic age estimation systems to help them estimate customers' ages for age-restricted products.
- Financial institutions could need AI-driven cybersecurity software to protect their computer systems.
- E-commerce online platforms may use AI-powered recommender systems to suggest products that their customers might like.

Since developing AS in-house can - in some cases - involve high costs and because many organisations do not have in-house skills to be able to build and deliver AI, many organisations choose to procure AS from third parties. Choosing to procure AS can lower the barriers to entry to state-of-the-art capabilities. In addition, there is some evidence to suggest that, in coming years, more organisations will turn to third parties for their AS instead of developing in-house. While in 2021, the global AI-as-a-service (AlaaS) market size was valued at \$5.6 billion (£4.5 billion), it is predicted to expand at a compound annual growth rate of 37.1% from 2022 to 2030.⁴

While there are advantages to procuring AS, there are risks too. In our research as part of the DRCF paper '*The benefits and harms of algorithms*'⁵ we heard concerns from academics about the lack of transparency and accountability during the procurement process. Research into business-to-business AI services has found that the roles of "controller", "joint controller" and "processor" - as defined in data protection legislation – are not always clearly identified, which gives rise to potential concerns that those building, selling and using AS may not be fulfilling their obligations under the UK GDPR.⁶

² Digital Regulation Cooperation Forum (2022) The benefits and harms of algorithms: A shared perspective from the four digital regulators. Available: <https://www.gov.uk/government/publications/findings-from-the-drcf-algorithmic-processing-workstream-spring-2022/the-benefits-and-harms-of-algorithms-a-shared-perspective-from-the-four-digital-regulators>

³ Digital Regulation Cooperation Forum (2022) Auditing algorithms: The existing landscape, role of regulators and future outlook. Available: <https://www.gov.uk/government/publications/findings-from-the-drcf-algorithmic-processing-workstream-spring-2022/auditing-algorithms-the-existing-landscape-role-of-regulators-and-future-outlook>

⁴ Grand View Research. Artificial Intelligence as a Service. Available: <https://www.grandviewresearch.com/industry-analysis/artificial-intelligence-as-a-service-market-report>

⁵ Digital Regulation Cooperation Forum (2022) The benefits and harms of algorithms: A shared perspective from the four digital regulators. Available: <https://www.gov.uk/government/publications/findings-from-the-drcf-algorithmic-processing-workstream-spring-2022/the-benefits-and-harms-of-algorithms-a-shared-perspective-from-the-four-digital-regulators>

⁶ Cobbe, J. and Singh, J. (2021). Artificial Intelligence as a Service: Legal Responsibilities, Liabilities, and Policy Challenges. Computer Law & Security Review. Available: <https://www.sciencedirect.com/science/article/abs/pii/S0267364921000467>

We also identified the following issues which can lead to a lack of transparency in the AS procurement process:

- Insufficient information from vendors to support organisations to understand the benefits of the AS or how they can mitigate risks.
- Low data maturity⁷ and skills among buyer organisations which means that buyers may not: understand their needs for an AS, use the AS for optimal effectiveness, or effectively scrutinise its quality and functionality.
- Few standardised approaches in place within vendor or buyer organisations to guide the procurement process.
- Insufficient information from vendors on how legal or ethical requirements have been met during the development of the AS ahead of deployment, for example, the representativeness and relevance of the training data. This can inhibit buyers from making an informed judgement about whether the AS is appropriate for their context.
- The commercial confidentiality of vendor-specific information, which can hinder independent testing of algorithmic systems procured by a buyer.

In this paper, we have used the term **'vendor'** to describe an organisation that develops and/ or sells AS. We have used **'buyer'** to refer to an organisation that procures and deploys AS. However, we found that the ecosystem does not neatly map onto binary categories: vendors may be the intermediary between an initial developer and a buyer, while buyers may also develop their own AS in-house.

In this paper, we interpret **'transparency'** broadly and include the establishment of clear lines of accountability between the vendor and buyer; the mutual sharing of valuable information about the AS; and the ability to explain and interpret how an AS has made a specific decision.

Finally, we refer to **'procurement'** activities as including both 'pre-procurement' activities (i.e., what buyers do before they procure an AS) and 'post-procurement' activities (i.e., what services are provided by vendors after agreements and contracts have been signed).

⁷ Roughly, whether an organisation can effectively use data to its fullest extent. See: <https://www.dataorchard.org.uk/what-is-data-maturity>; <https://www.gov.uk/government/news/introducing-the-government-data-maturity-model>

Findings from our workshops on transparency in the procurement of algorithmic systems

Finding # 1: There is not a one-size-fits-all approach to achieving transparency in the procurement of algorithmic systems

In the workshops, we asked participants about the procurement process for AS, who is involved from vendor and buyer organisations, and how and what types of information is shared.

Understanding what information to provide and request will vary depending on what type of algorithmic tools are procured

We found that the procurement of algorithmic systems can take several different forms, including:

- A buyer procuring an 'off-the-shelf' system (i.e., a system that is fully developed and ready to be deployed).
- A buyer providing inputs into a vendor's model which could produce an output for the buyer (sometimes referred to as AI-as-a-service).
- A vendor providing tools that enable buyers to develop their own algorithmic systems (sometimes referred to as 'no code' or 'low code' solutions). These tools can allow buyers with limited coding expertise to develop their own AS.
- A vendor providing data sets to buyers, which can be used to train their own AS.
- A buyer using open-source, free-to-use AS.

Organisations have different processes for procuring AS

However, amongst these, we identified several common stages across the procurement lifecycle. These include:

- Buyers identifying the requirement for an algorithmic tool.
- Buyers identifying the need for procurement of an AS.
- Buyers engaging with a vendor's sales team to enquire about their products.
- Buyers completing due diligence of the AS and the vendor. For example, seeking technical, governance and compliance-related information from vendors.
- Vendors completing tests to consider 'market fit'. For example, assessing how well their AS is likely to work in the buyer's context.
- Vendors providing training and support for those who will use the AS, such as buyers or third parties.
- Buyers and vendors completing activities to monitor, review and update the AS post-deployment.
- Buyers and vendors engaging in further procurements or ending contracts and service agreements with one another.

The information exchanged between vendors and buyers prior to procurement can vary in quantity, detail, and utility (for the buyer).

Workshop participants told us the types of information that they would usually share and request about an AS before procurement takes place. We observed that the range of information shared, and the detail of that information, varied between organisations.

The types of information shared by vendors and buyers can be broadly divided into two categories:

1. Information about how the AS performs and should be used.
2. Information about how the AS was designed and developed.

Category 1 examples: Information about how the AS performs and should be used	Category 2 examples: Information about how the AS was designed and developed
<ul style="list-style-type: none"> • The capabilities and limitations of the AS. • Technical features of the AS. • Information about procedures and protocols that govern its use. • Information about how buyers can facilitate requests for redress by citizens affected by the outputs of the AS. • Inputs that buyers should put into the AS. • Types of outputs it generates. • How to interpret the outputs of the AS. • How to use the AS effectively and appropriately, e.g., whether it needs to be overseen by a human operator. 	<ul style="list-style-type: none"> • Any impact assessments undertaken (e.g., data protection impact assessments). • Where data was sourced to train the AS. • What data was used for training the AS. • How AS were tested. • What performance metrics were used for testing. • What ongoing support will vendors provide for the AS (e.g., whether vendors will continue to develop the AS, how they will do this, and how the updates will be cascaded down to buyers).

Workshop participants shared several methods of engaging with the other party to share information. For example, where a general, off-the-shelf system is being procured, vendors told us that they would tend to share information through generic product marketing material, technical documents, reports and APIs,⁸ public policy papers or user guides. However, in a bespoke development or adaptation of an AS, vendors told us that they may also set up consultations, calls, and sometimes provided additional services like training or advice.

⁸ API is Application Programming Interface: a way for computer programs to communicate with each other. See more: <https://www.redhat.com/en/topics/api/what-are-application-programming-interfaces>

Assessments of whether AS will be deployed appropriately and lawfully are undertaken by both vendors and buyers

Due diligence is a process that typically involves evaluating the assets or liabilities of a company, its goods, or services. In our workshops, we heard examples of both vendors and buyers undertaking due diligence to help them make sure that the AS would be deployed appropriately.

Participants of the workshop referred to the following examples of due diligence:

- Vendors may complete a product-market fit evaluation, to help ensure that their AS met the buyer's criteria and deployment context.
- Vendors may employ risk assessment methods to check whether the buyer may deploy its AS in an unfavourable use case or a way that the vendor does not intend for.
- Vendors and buyers may check whether the use of the AS would comply with relevant legal requirements in the country of deployment. For example, checking whether an AS developed in the United States would be compliant if it was deployed in the UK. Buyers may also consider an AS's compliance with their own legal requirements.
- Vendors and buyers may specify 'red line' industries, markets, and countries where an AS should not be deployed.
- Vendors may provide and buyers may seek assurance that the AS works effectively in the country of deployment (e.g., some AS may struggle to recognise non-European languages, character sets or local dialects).
- Vendors and buyers may establish governance processes to oversee decisions relating to procurement, for example via an internal or external ethics group.
- Vendors may build levers into a product or agreement to stop abuse of an AS once procured, and to 'turn off' the service provided to a buyer.
- Buyers may undertake a technical interrogation of the AS (e.g., via a demo or technical documentation).
- Buyers may consider the vendor's client base (e.g., buyers may check to see whether the AS is used by other businesses in their sector).
- Vendors and buyers may consider one another's stated values or internal standards.
- Buyers may develop a business case to analyse and determine the likelihood of the effectiveness of the AS being procured.

A variety of support (including no support) may be provided after procurement has taken place

There is likely to be a need for vendors and buyers to maintain a relationship after the initial procurement of an AS. This can include training and support for end-users, as well as monitoring, reviewing, and updating the AS. Often, vendors and buyers will agree on the type of engagement and support that is needed after the tool has been procured, when agreeing the terms of procurement.

Human training and support for end users was not a widespread practice among participants of our workshop. Instead, we heard that vendors more often provide several types of arms-length support including:

- Publishing updated technical documentation and supporting material as new model versions are released.

- Publishing updates to a policy document which sets out the methods used to develop and train a model.
- Providing support at a distance (e.g., via online collaboration and communication tools like Slack).
- Providing services as an API or a Software Development Kit (SDK).⁹

Bespoke staff training and support after procurement is likely to become more important for companies who do not have appropriate in-house technical expertise.

Participants also gave a mix of responses to a question on the extent to which vendors monitored and updated an AS being deployed by a buyer. Where the system uses machine learning (ML), there may be a need for regular monitoring of the performance of an AS to avoid risks such as ‘drift.’ Drift in the context of ML is the decay of a model’s predictive ability over time. This can occur for several reasons such as data drift (where properties of the underlying data changes) or concept drift (where the properties of what is being predicted change).

Finding #2: Buyers can lack the technical expertise to effectively scrutinise the AS they are procuring, whilst vendors may limit the information they share with buyers

In the workshops, we wanted to better understand the barriers that can inhibit transparent and accountable information-sharing in the procurement process. While there may be several reasons for this, three stood out:

Some buyers can lack the necessary expertise to scrutinise the AS they wish to procure

Participants at our workshops told us that it is typical for vendors’ sales teams and buyers’ procurement teams to manage the procurement process and may work with a project or business team to explore their needs. Buyers need to know that the AS being procured meets their needs and has been developed and tested responsibly. However, it requires a certain level of technical expertise to understand key details about an AS and to be informed consumers.

In our workshops, we heard that some buyers lacked understanding of AS and could struggle to recognise where an algorithmic process had been integrated into a system they were procuring. For example, we heard that procurement teams had asked vendors simple questions, for example: “Does this solution contain an algorithm?” This issue may be compounded where vendors fail to note that a solution includes AI or its subset, ML. This can also be difficult for buyer organisations to understand and oversee what AS have been procured, or where and how they are being applied across the business.

We also heard that a lack of expertise among buyers could have a disproportionate impact on small- and medium- sized vendors, who may need to spend time responding to questions. Vendors told us that they had received template procurement questionnaires from buyers, which included a significant number of questions that vendors did not believe were relevant to the AS.

⁹ Software Development Kit is a collection of tools and programs developers can use to build on specific software or hardware platforms. See: <https://www.techtarget.com/whatis/definition/software-developers-kit-SDK>

Where buyers have insufficient information about the development or testing of an AS, there is a risk that buyers could be deploying an AS that is unlawful or unethical. This risk is particularly acute for high-risk applications of AS, for example where an AS determines a person's access to employment or housing or where the application is in a highly regulated sector such as finance.

Vendors may cherry pick performance metrics which can over-state the effectiveness of an AS

Our workshops found that vendors use a range of performance metrics and testing methods. However, without appropriate technical expertise or scrutiny, these metrics may give buyers an incomplete picture of the effectiveness of an AS.

Some participants suggested that vendors share performance metrics that overstate the effectiveness of their AS, whilst omitting other metrics which indicate lower effectiveness in other areas. Some vendors raised concerns that their competitors choose the most favourable (i.e., the highest) performance metric to win procurement contracts. While we heard that some buyers do request specific performance metrics, not all buyers may have the technical knowledge to understand which performance metrics are most relevant to their procurement decision.

Alternative methods of distorting algorithmic effectiveness also exist. These include: choosing a test dataset that will portray an AS in a more favourable light (e.g., age verification being tested on pictures taken under ideal conditions); or ignoring subsets of the population that an algorithm has trouble classifying.

Some vendors also felt frustrated about what they saw as the current market not prioritising transparency in the testing process. We heard concerns that without requirements for independent testing and validation of performance, there are few incentives for vendors to do this. Some vendors felt that this meant competitors could cite better 'performance' figures that were not validated or independently assured.

Other important metrics - including algorithmic fairness, lack of bias, and resilience to external changes - were not discussed at length during the workshops. Similar challenges exist in developing appropriate definitions for these areas, which can be contradictory and context-dependent and will require technical expertise and input from organisations.¹⁰ For example, there are a broad range of definitions for algorithmic fairness, and it is important that these can be understood and compared by buyers to decide which is most useful for their context.

Vendors can find it difficult to strike the right balance between providing full information about AS and protecting commercially sensitive information

In the workshops, vendors told us that they remain concerned about revealing commercially sensitive information to buyers. Some vendors felt that sharing too much technical detail or knowledge could allow buyers to re-develop their product. However, vendors also recognised that their desire to protect commercial interests should be balanced with the need to provide buyers with sufficient information to be confident in the AS they are procuring.

¹⁰ Centre for Data Ethics and Innovation (2021) Review into bias in algorithmic decision-making. Available; https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/957259/Review_into_bias_in_algorithmic_decision-making.pdf

While the exchange of information between firms can generate efficiencies and benefits, certain exchanges of information – notably those of commercially sensitive information with competitors – may reduce strategic uncertainty in the market to such an extent that the exchange is restrictive or distortive of competition. Before taking a ‘share-all’ approach, firms should make sure that any sharing of information would not breach the provisions of the UK Competition Act 1998 (or indeed the competition law of any other relevant jurisdiction).¹¹

Finding #3: We are already seeing some promising solutions that help to improve transparency

During the workshops, we explored existing and emerging solutions to address barriers to transparency. While several solutions were discussed, below we have set our own views on areas that we think would be valuable to explore in more detail.

Certification, standards, and guidelines may have a role in helping to set out agreed metrics that vendors can share with buyers

Certification, standards, and guidelines have the potential to:

- Help businesses to follow agreed good practice, which may (although not always) also help them to comply with regulation.
- Ensure that businesses are competing with one another on an equal footing and are protected from others acting unfairly.
- Provide internationally agreed definitions, requirements, measures, and evaluation processes.

In the workshops, we heard that common or standard metrics do not yet exist within industry for the evaluation of AS. For vendors, this can make it more challenging to provide useful information, and for buyers, this lack of consistency can make it difficult to compare different AS.

Buyers also told us that they would find more detail on the performance of the AS being procured helpful - including across a range of metrics. The development of more consistent performance metrics could also help regulators to better understand how accurate an AS is in a specific context.

In the workshops, vendors and buyers appeared to agree that some form of standardisation on the minimum information that should be made available to buyers would be useful. However, we also heard concerns that many buyers lack technical knowledge to understand the implications of some metrics; and that this could create burdens for smaller vendors if they were required to complete lengthy and time-consuming standards at each procurement.

In our desk research, we identified several examples of certification, standards and guidance that may be helpful for both vendors and buyers of algorithmic systems, as set out in the remainder of the section below.

The Institute of Electrical and Electronics Engineers is developing a **standard for transparency in autonomous systems**. This may aid vendors in providing consistent and clear metrics and buyers in comparing solutions presented by vendors. The standard sets out measurable, testable levels of

¹¹ Competition and Markets Authority (2014) How to manage competitively sensitive information. Available: <https://www.gov.uk/government/publications/limiting-risk-in-relation-to-competitors-information>

transparency for autonomous systems, so that they can be objectively assessed, and compliance determined.¹²

In addition, the British Standards Institute has published a Publicly Available Specification for online age verification.¹³ This intends to assist vendors of age restricted products and services online with a means to demonstrate best practice and compliance regarding age checking and provide assurance for buyers

- Age check policies.
- Approach to quality assurance and internal audit.
- Technical programming and performance.

While there are multiple guidelines for public sector procurement of AI¹⁴ we found fewer guidelines for private sector procurement of AS. One guide which may be helpful is the ICO's jointly produced guidelines with the Alan Turing Institute on Explaining Decisions Made With AI.¹⁵ This guidance aims to give teams within organisations - including compliance, technical and senior management teams - practical advice to help explain the processes, services and decisions delivered or assisted by AI, to the individuals affected by them.

Templates and registers can function as a guide to help organisations understand good practice around types of information to share

There may be a role for algorithmic templates and registers to help guide vendors to share and present information clearly. Templates can help to organise information about AS in a structured way that can be aimed at both technical and non-technical experts.

The UK government has developed an algorithmic transparency recording standard which provides a template to help public sector organisations share information about how they develop and use AS. It includes several categories of information that are required to be shared, including describing the datasets the model has been trained on, a link to a data protection impact assessment, and a list of common risks in using the tool (for example, using it in a way that it was not intended for).¹⁶ The government has also created an algorithmic transparency recording standard hub which includes a

¹² 7001-2021 IEEE Draft Standard for Transparency of Autonomous Systems. Available: <https://ieeexplore.ieee.org/document/9451892>

¹³ Age Verification Providers Association. International standard for age verification. Available: <https://avpassociation.com/standards-for-age-verification/>

¹⁴ UK Government (2020) Guidelines for AI procurement. Available: <https://www.gov.uk/government/publications/guidelines-for-ai-procurement/guidelines-for-ai-procurement#top-10-considerations>; Ada Lovelace Institute, AI Now Institute and Open Government Partnership (2021) Algorithmic accountability for the public sector. Available: <https://www.opengovpartnership.org/wp-content/uploads/2021/08/algorithmic-accountability-public-sector.pdf>; Ada Lovelace Institute (2020) Transparency mechanisms for UK public-sector algorithmic decision-making systems. Available: <https://www.adalovelaceinstitute.org/wp-content/uploads/2020/10/Transparency-mechanisms-explainer-1.pdf>; Centre for Data Ethics and Innovation (2020) Review into bias in algorithmic decision-making. Available: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/957259/Review_into_bias_in_algorithmic_decision-making.pdf

¹⁵ Information Commissioner's Office, Alan Turing Institute (2020) Explaining decisions made with AI. Available: <https://ico.org.uk/for-organisations/guide-to-data-protection/key-dp-themes/explaining-decisions-made-with-ai/>

¹⁶ Central Data and Digital Office, Centre for Data Ethics and Innovation (2021) Algorithmic transparency recording standard. Available: <https://www.gov.uk/government/publications/algorithmic-transparency-template>

collection of published transparency reports.¹⁷ There may be value in vendors and buyers drawing on categories of information in these templates as a guide for sharing information.

Another example is that vendors may publish documentation detailing the performance characteristics of models, sometimes referred to as model cards.¹⁸ Model cards can provide benchmarked evaluation in different contexts. For example, a model card could show how well a speech-to-text translation model works in a health setting in comparison to an environment where many speakers have English as a second language. Through model cards, vendors can also disclose where the models should and should not be deployed, and other useful information.

An AI auditing market may help to provide assurance to buyers that a technology has met a certain threshold for quality

In the AI audit market, firms offer a commercially driven service to provide independent assessment of an algorithmic system against a framework. In our workshops, we heard calls from several vendors and buyers to explore the idea of encouraging independent audits of AS products. They felt audits could help to establish more consistent metrics to ‘measure’ products against, and that this in turn could help create a level playing field for vendors and enable buyers to more easily compare solutions.

There are a number of questions around the role of third-party auditors and how they would operate. For example, who would pay for the audit? Would a vendor or buyer require them? Who would be required to be audited?

An additional consideration is that independent audits may incur a financial cost for vendors, which may disproportionately burden smaller vendors. In addition, vendors could be disincentivised from seeking independent audit where other companies aren’t doing the same.

The DRCF algorithmic processing team has a separate workstream that is undertaking an initial mapping exercise of the AI auditing market, so that we can understand how these services are developing.

¹⁷ Central Digital and Data Office, Centre for Data Ethics and Innovation (2023) Algorithmic transparency recording standard hub. Available: <https://www.gov.uk/government/collections/algorithmic-transparency-recording-standard-hub>

¹⁸ Model cards have been developed and used by several companies including Google, Microsoft, Salesforce, Nvidia, OpenAI, Meta, DeepMind. See Github for a list of Model Cards and DataSheets: <https://github.com/ivylee/model-cards-and-datasheets>

Key legislative and international developments

There are several live legislative and international developments that will likely affect discussions in this space.

The UK Government's AI policy paper and forthcoming AI White Paper

In July 2022, the UK Government published a policy paper and public consultation outlining proposals to regulate AI. In it, the Government sets out a series of non-statutory AI principles, including making sure that AI is appropriately transparent and explainable, and clarifying routes to redress or contestability. The Government is due to publish a White Paper on AI in Spring 2023.

The draft European Commission's AI Act (EU AI Act)

The draft EU AI Act proposes specific provisions around transparency and what information is provided to users. While still being developed, the Act will likely apply to those who sell AI systems in the European Union. When enacted, it could also have implications for the development of international standards.

The European Commission's AI Liability Directive

This Directive proposes to develop uniform rules for certain aspects of non-contractual civil liability for damage caused with the involvement of AI systems. If enacted, the Directive may have implications for vendors and buyers of algorithmic systems who are subject to it. For example, the European Commission has proposed a right to evidence about how AI systems, which have allegedly led to injury, or an abuse of fundamental rights, were designed, developed and deployed.

Conclusion

The findings set out in this paper will help to inform future research and policy by individual regulators within the DRCF, including on algorithmic age assurance, content moderation, algorithmic assessments and audit.

The DRCF algorithmic processing team intends to continue to explore key themes raised in this research in its 2023-2024 work programme.

More broadly, transparency in the context of AS will remain a key focus within the UK, the European Union and internationally as jurisdictions develop their approaches to AI regulation.